Notes

* This paper has greatly benefited from the advice and comments of Bruce Hayes and Donca Steriade. I am also grateful to Robert Kirchner and Jie Zhang for their helpful comments.
1. See Jun 1995 for the discussion of the typological data and their original sources.
2. In Catalan, among continuants, only /l/ may assimilate in place to a following palatal or velar but it cannot to a following labial (Mascaró 1978). This limited case is not specified in (8) for simplicity's sake.
3. The assimilation of labials in the relevant dialects is a well-established diachronic fact; however, because of the lack of labial-final stems (Bobaljik 1996: 323), it is harder to show synchronically. Preconsonantal labials are phonologically illegal except before another labial, a gap that could easily be accounted for by assuming assimilation.
4. Broad phonetic transcriptions are employed for these examples. For instance, actual phonetic forms are subject to the regular process of post-obstruent fortition, whereby lenis obstruents become fortis after an obstruent (Kim-Renaud 1986).
5. See section 5 below for gradient assimilation, in which the loss can be partial.
6. Kirchner (1998, this volume) and Boersma (1998) adopt the same type of constraint.
7. In assessing violations of the WEAKENING, short and long gestures are not differentiated.
8. For the analysis shown in (15), a perceptual definition of faithfulness provided in (14) does not seem to be necessary; a conventional articulatory definition will be sufficient. However, in section 5, the role of the perceptual faithfulness will be evident in the analysis of gradient place assimilation.
9. If the gap is genuine, it cannot be explained by a universal ranking PRESERVE(pl(__V)) ≫ WEAKENING, since this ranking would exclude all cases of progressive assimilation in suffixes (section 3.3).

4   The typology of rounding harmony

Abigail R. Kaun

1   Rounding harmony typology

Rounding harmony is a phonological process whereby certain vowels surface as rounded under the influence of a neighbouring rounded vowel. What is striking about rounding harmony is the fact that the simplest possible statement – 'a vowel must be rounded when preceded by/followed by a rounded vowel' – fails to characterise the great majority of rounding harmony systems. In most cases, conditions referring to tongue body position (height and/or backness) are imposed on either the triggering element, the target, or both. I argue that this interaction among vowel features renders traditional rule-based accounts of the typological patterns nonexplanatory. Within a constraint-based framework such as Optimality Theory (Prince and Smolensky 1993), however, the interaction of rounding with these other phonological dimensions can be modelled in a straightforward manner that allows for the characterisation of all attested rounding harmony patterns, while making falsifiable predictions regarding the logically possible but cross-linguistically unattested patterns. Central to my analysis is the claim that phonological systems are organised around principles of articulation and perception. These principles are encoded in the formal grammar as Optimality-Theoretic constraints.

The goals of the chapter are as follows: (1) to exemplify the range of attested rounding harmony patterns, (2) to identify the perceptual and articulatory principles that give rise to these patterns, and (3) to propose a formal model to characterise the role of these principles in grammar. In the final section, I outline the results of a recent experiment involving loanwords in Turkish. The experimental results indicate that the model of rounding harmony developed here, while motivated by evidence from typology, is also an appropriate model of individual grammars.

For data, I rely on a number of earlier typological studies of rounding harmony (Bogoroditskij 1953; Korn 1969; Kaun 1994, 1995), as well as sources describing rounding harmony in Mongolian (Svantesson 1985), Tungusic (especially Li 1996), and a variety of non-Altaic languages. In total, thirty-three languages were surveyed.

Table 4.1. *Height-conditioned typology*

Type	Same-height harmony		Cross-height harmony		Sample language, number of languages
	high target	nonhigh target	high target	nonhigh target	
1	✓	✓	✓	✓	Kirgiz, 2
2	✓	✓			Yokuts, 1
3	✓				Hixkaryana, 3
4		✓			Khalkha Mongolian, 8
5	✓		✓		Nawuri, 7
6	✓	✓	✓		Yakut, 2

What all of these languages have in common is that rounding harmony is nearly always constrained so as to apply only when certain conditions are met: conditions that refer to phonological dimensions other than lip rounding.

1.1   Rounding harmony and height

Both vowel height and backness play a role in rounding harmony. I will begin by addressing harmony systems in which it is only the height of the trigger and/or target that determines whether harmony will apply. Table 4.1 gives the typology of height-sensitive systems. For each type, a representative language is listed, followed by the number of languages in the sample instantiating that particular pattern. A check mark indicates that rounding harmony is observed in the relevant configuration. In 'same-height' harmony the trigger and target agree in height, whereas in 'cross-height' harmony the trigger and target disagree in height.

In the Type 1 system, rounding harmony is triggered by any rounded vowel and targets any vowel. In the majority of attested rounding harmony systems, however, harmony is constrained by vowel height. Same-height harmony is more frequently observed than cross-height harmony (nine types vs four), and where cross-height harmony arises, it typically involves a nonhigh trigger and a high target, that is, *oCu/öCü* sequences rather than *uCo/üCö* sequences, where the first vowel represents the trigger. From this, we may conclude that the effects of height on rounding harmony are limited to the adoption (within particular languages) of one or more of the following principles:

(1) i. The trigger must be nonhigh.
    ii. The target must be high.
    iii. The trigger and target must agree in height.

In the Yokuts type (Type 2), rounding harmony occurs as long as condition (iii) is met, that is, as long as the trigger and target agree in height. In the Type 3

languages (Hixkaryana, Kachin Khakass, and Tsou), both conditions (ii) and (iii) must be met in order for harmony to apply; hence only high vowels trigger harmony, and only high vowels undergo it. Type 4, widely observed in the Mongolian and Tungusic languages families, requires both (i) and (iii), so that harmony is observed only among nonhigh vowels. Harmony in Type 5 applies as long as (ii) is met. Finally, harmony applies in the Type 6 system as long as either (ii) or (iii) is met, so that the only configuration where harmony is blocked is where the trigger is high and the target is nonhigh, ruling out the widely dispreferred *uCo/üCö* sequences.

Let us examine data from two of these types by way of illustration. The most general pattern, that which imposes none of the height restrictions from (3), is Type 1. The dialect of Kirgiz reported in Comrie (1981) instantiates this system. Kirgiz has a fully symmetrical vowel inventory of sixteen vowels, classified as front/back, rounded/nonrounded, high/nonhigh, and long/short; thus /i, iː, ü, üː, ɪ, ɪː, u, uː, e, eː, ö, öː, a, aː o, oː/. The quality of vowels in noninitial syllables is to a large extent predictable on the basis of the quality of the vowel occurring in the first syllable: vowels agree with the vowel of the initial syllable in both backness and rounding. The effects of backness harmony and rounding harmony are most readily apparent in suffixal vowel alternations, although the vowels of native polysyllabic roots display the same distributional patterns.

Let us consider first the ordinative suffix, which has the surface variants {*-(i)ntʃi, -(ɪ)ntʃɪ, -(u)ntʃu, -(ü)ntʃü*}. The vowels of this suffix are in all instances high; however, their rounding and backness is variable. When the root contains front unrounded vowels the alternant *-(i)ntʃi* surfaces. Following back unrounded vowels the suffix contains back unrounded vowels and the alternant *-(ɪ)ntʃɪ* surfaces: thus *bir* 'one', *bir-intʃi* 'one-ord.'; *beʃ ~ beʃ-intʃi* 'five'; *altɪ ~ altɪ-ntʃɪ* 'six'; *ʒɪjɪrmɪ ~ ʒɪjɪrmɪ-ntʃɪ* 'twenty'. The vowels of this suffix are rounded following roots containing rounded vowels, as in the following examples: *ütʃ ~ ütʃ-üntʃü* 'three', *tört ~ tört-üntʃü* 'four', *toguz ~ toguz-untʃu* 'nine', *on ~ on-untʃu* 'ten'.

To demonstrate the effects of backness and rounding harmony in nonhigh vowels, consider the ablative suffix, which has the surface variants {*-t/den, -t/dan, -t/dön, -t/don*}. The nonhigh vowel of this suffix also agrees in both backness and rounding with the vowels of the root. (Additionally, consonants agree in voicing with a preceding consonant): *iʃ* 'work' ~ *iʃ-ten* 'work-ABL', *et ~ et-ten* 'meat', *ʒɪl ~ ʒɪl-dan* 'year', *alma ~ alma-dan* 'apple', *üj ~ üj-dön* 'house', *köl ~ köl-dön* 'lake', *tuz ~ tuz-don* 'salt', *tokoj ~ tokoj-don* 'forest'. This effect is pervasive across sequences of suffixes, as illustrated in polymorphemic words such as *köz-ün-dö* 'eye-poss.-loc.'

This pattern, while simple and symmetric, is in fact very unusual. More typical is the familiar pattern instantiated in Turkish (Type 5), where vowel height plays a role in determining the application of harmony. The vowels of

Table 4.2. *Backness-sensitive typology (partial)*

| Type | Front trigger | Back trigger | | | | Sample language, number of languages |
| | | Same-height | | Cross-height | | |
		high target	nonhigh target	high target	nonhigh target	
7	✓	✓	✓	✓		Altai, 2
8	✓	✓		✓		Karakalpak, 3
9	✓	✓				Kyzyl Khakass, 1

Turkish, like those of Kirgiz, contrast for frontness, rounding and height: /i, ü, ı, u, e, ö, a, o/. Also like Kirgiz, suffix vowels undergo both backness and rounding harmony. While backness harmony functions entirely independently of vowel height, height plays a role in rounding harmony. High suffixes, such as the first singular possessive, undergoes rounding harmony, as in *ip* 'rope', *ip-im* 'rope-1.sg.poss.', *süt ~ süt-üm* 'milk', *ev ~ ev-im* 'house, *tʃöp ~ tʃöp-üm* 'garbage', *kız ~ kız-ım* 'girl', *buz ~ buz-um* 'ice', *at ~ at-ım* 'horse', *gol ~ gol-um* '(football) goal'.

Nonhigh vowel suffixes, such as the dative suffix, do not undergo rounding harmony: *ip* 'rope' ~ *ip-e* 'rope-dat.', *süt ~ süt-e* (*süt-ö*) 'milk-', *ev ~ ev-e* 'house', *tʃöp* 'garbage' ~ *tʃöp-e* (*tʃöp-ö*) 'garbage', *kız ~ kız-a* 'girl', *buz ~ buz-a* (*buz-o*) 'ice', *at ~ at-a* 'horse', *gol ~ gol-a* (*gol-o*) '(football) goal'.

Data from languages exhibiting the remaining height-sensitive types are presented in Kaun 1995.

1.2   Rounding harmony and backness

In addition to vowel height, backness also emerges as a conditioning factor in the typology of rounding harmony. In a number of Turkic languages, height conditions of the sort just described are imposed when the trigger is a back vowel, but are suspended when the trigger is front, thus yielding across-the-board harmony among front vowels. Adding these cases to our typology yields three additional types, given in table 4.2. Check marks in the first column indicate that when the trigger is a front vowel, no height restrictions are imposed. Each of these types can be thought of as a version of one of the height-sensitive types from table 4.1. For example, Type 7 can be thought of as a variant of Type 6 with the stipulation that all height conditions are suspended when the trigger is front. Similarly, Type 8 is a variant of Type 2, in which harmony triggered by back vowels may only target high vowels. Type 9 can be thought of as a variant

Table 4.3. *Backness-sensitive typology (continued): back targets preferred*

| Type | Front trigger | Back trigger | | | | Sample language, number of languages |
| | | Same-height | | Cross-height | | |
		high target	nonhigh target	high target	nonhigh target	
10	(*)	✓	✓	✓		Sibe (roots), 1
11	(*)	✓		✓		Sibe (affixes), 1
12	(*)		✓			Shuluun Höh, 1

of Type 3 in that the operative height constraints are the requirement that the trigger and target agree in height, and that the target must be high.

In Shuluun Höh (Svantesson 1985), like other Mongolian languages, rounding harmony is observed as long as the trigger and target agree in height and the trigger is nonhigh. Shuluun Höh imposes the additional requirement that the target of harmony must be [+back]. Rounding harmony in Sibe, a Tungusic language of China (Li 1996), imposes this condition as well, and will be discussed in greater detail in section 3. Shuluun Höh and Sibe provide three additional types (Sibe itself exhibiting two: one within roots and one across morpheme boundaries).

Through extensive fieldwork on Turkic languages in Siberia, Harrison (2001) has uncovered several rounding harmony systems that appear to be in transition. In Tofa, the trigger of rounding harmony must be a front vowel, and the target must be high. Some speakers exhibit harmony consistently, others sporadically, and still others not at all. The same type of variability is evidenced in Tuha, where the target of rounding harmony is always high, and harmony is consistently applied when the trigger and target agree in height, yielding *uCu* and *üCü* sequences. When the trigger and target are distinct in height, harmony is applied only variably. Thus in Tuha, one hears harmonic *oCu/öCü* sequences as well as nonharmonic *oCı/öCi* sequences. In Altai Tuvan, rounding harmony may occur in any trigger/target pair, but is apparently obligatory only when the trigger is the nonhigh front vowel [ö].

I will treat Altai Tuvan as essentially an exemplar of Type 1, in which harmony applies without regard to the backness or height of the trigger and target, though taking note of this case as an additional instantiation of the preference for nonhigh and front triggers, as it is only the nonhigh front vowel [ö] that obligatorily triggers harmony. Tuha can be thought of as an exemplar of Type 5, exemplified most familiarly in Standard Turkish; however, this language provides an additional case in which same-height harmony is preferred over cross-height harmony. Tofa constitutes a new type (labelled Type 13 in

Table 4.4. *Backness-sensitive typology (completed)*

Type	Back trigger	Front trigger				Sample language, number of languages
		Same-height		Cross-height		
		high target	nonhigh target	high target	nonhigh target	
13		✓		✓		Tofa, 1

table 4.4), in which only front vowels trigger harmony and only high vowels are available as rounding harmony targets:

1.3   Summary

To summarise, all of the observed types can be characterised by means of the five height and backness conditions given in (2):

(2) Conditions favouring rounding harmony
    i.   The trigger is nonhigh.
    ii. The trigger is front.
    iii. The target is high.
    iv. The target is back.
    v. The trigger and target agree in height.

From this list we should take note of the fact that the preferred rounding harmony *targets*, namely the high vowels and the back vowels, are the typologically dispreferred rounding harmony *triggers*. I will show that this semi-complementary distribution is not an accident.

In the remainder of this chapter, I discuss the phonetic underpinnings of this typology and present a formal account that incorporates these general phonetic principles as its foundation. Section 2 lays out the phonetic properties of rounded vowels, the key observation being that when lip rounding is combined with diverse tongue shapes, its articulatory, acoustic, and perceptual manifestations are nonuniform. In this section, I also propose that vowel harmony is perceptually driven and that perceptual differences associated with vowels of differing tongue shapes are the source of the observed trigger/target asymmetries. In section 3, I posit Optimality-Theoretic constraints to account for the typology and demonstrate that the proposed analysis provides a good fit with typological data. In the final section, I describe the results of an experiment on harmony in loanwords in Turkish, and suggest that those experimental results provide support for the analysis not only as a metagrammar for the typology as a whole, but also as a model of individual grammars.

2   Phonetic underpinnings of the rounding harmony typology

We turn now to the question of phonetic grounding. The crucial questions are: Why should nonhigh vowels be preferred as triggers, and high vowels as targets? Why is harmony preferred when the trigger and target agree in height? And, more generally, why is rounding harmony sensitive to vowel features other than rounding, namely to height and backness? In seeking an answer to these questions, I will first discuss the articulatory, acoustic, and perceptual properties of rounded vowels.

2.1   Articulation

Linker (1982) studied labial activity in vowels for five genetically diverse languages: English, Cantonese, Finnish, French, and Swedish, with the goal of identifying the linguistically significant parameters of lip position. With the exception of English, both back and front rounded vowels are present in the inventories of all of the languages studied, and all of the languages exhibit both high and nonhigh rounded vowels. Thus, the data set allowed for the comparison of labial activity in back versus front rounded vowels, as well as a comparison of labial activity in high versus nonhigh rounded vowels.

Linker's study included measurements of twenty-four distinct dimensions taken from still photographs of the side and front view of the mouth.[1] Using the factor analysis algorithm PARAFAC (Harshman 1970; Harshman, Ladefoged, and Goldstein 1977; Harshman and Berenbaum 1980), Linker identified the articulatory dimensions of lip position that appear to be relevant for distinguishing vowels within each of the languages studied. These dimensions involved horizontal opening, vertical opening, and lip protrusion, or some combination thereof. Additionally, using CANON (Goldstein n.d.), Linker isolated a set of canonical factors of lip position relevant to all of the languages studied.

Her results yielded two articulatory factors associated with lip rounding: one based largely on horizontal opening, and the second based on vertical opening and lip protrusion. While all of the rounded vowels that were studied clustered on the higher end of both scales relative to the unrounded vowels, there were systematic differences of degree among the rounded vowels themselves. For all languages studied, the high rounded vowels fell on the higher end of the scales relative to their nonhigh counterparts, thus the high vowels were in a sense 'more rounded' than the nonhigh vowels. The back vowels tended to fall on the higher end of the scales relative to their front counterparts, the notable exceptions being the nonhigh vowels of Cantonese and French. In those cases, the front [ö]-type vowel was slightly more protruded than the back [o]-type vowel. Aside from this exception, Linker's results indicate that the magnitude of lip rounding is relatively greater for high vowels than for nonhigh vowels and, though with less uniformity, for back vowels than for front vowels.

2.2   Acoustics

This generalisation is consistent with the acoustic patterning associated with rounded vowels. Stevens (1998: 294) makes the following observation: '[ɪ]increased prominence of the principal spectral peak, together with a lowered centre of gravity of the peak, can be considered as primary acoustic correlates of the rounding feature'. When rounded vowels are compared in terms of both of these acoustic properties, it is evident that the acoustic consequences of adding lip rounding to back vowels is greater than that achieved by adding lip rounding to front vowels. It is also evident that lip rounding has more dramatic acoustic consequences for high vowels than for mid vowels (see Stevens' figure 6.22 (1998: 293)).

2.3   Perception

The difference in lip activity associated with high vowels vs nonhigh vowels and back vowels vs front vowels (Linker), alongside the observed acoustic differences (Stevens), apparently gives rise to perceptual differences among rounded vowels as well. Terbeek (1977) presents an investigation of perceptual distances in the vowel space. In his study, Terbeek examined the perceptual distance among ten monophthongs {i, ü, e, ö, ɪ, u, a, o, æ, ɚ} with the goal of identifying the perceptual attributes according to which listeners perceive differences among vowels.

Speakers of English, German, Thai, Turkish, and Swedish served as subjects. For each of these subjects, some but not all of the monophthongs were similar to vowels occurring in the listener's native language. The data consisted of triadic comparisons of the test vowels in the context [bəb___] and the task was to determine which of the three stimuli sounded the most distinct from the others. From the responses collected, dissimilarity matrices were constructed, which then were submitted to a PARAFAC factor analysis algorithm (Harshman 1970; Harshman, Ladefoged, and Goldstein 1977; Harshman and Berenbaum 1980). Terbeek's PARAFAC analysis yielded a six-dimensional solution, indicating that six factors are relevant to the identification of vowels within a multidimensional space. These six dimensions correlate more or less with the standard phonological oppositions shown in (3):

(3)   Dimensions of vowel identification (Terbeek 1977)
Dimension 1:	back vs nonback (1)
Dimension 2:	back vs nonback (2)[2]
Dimension 3:	low vs nonlow
Dimension 4:	high vs nonhigh
Dimension 5:	round vs nonround
Dimension 6:	peripheral vs central

The results of Terbeek's investigation indicate that along the round vs nonround continuum, the rounded vowels are arranged as shown schematically in (4):

(4)   The round vs nonround continuum (Terbeek 1977)
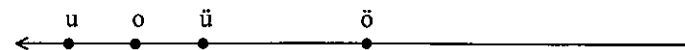
This arrangement indicates that nonhigh vowels and front vowels are perceived as relatively less rounded: the high vowels lie on the higher end of the scale relative to the nonhigh vowels, and the back vowels lie on the higher end of the scale relative to the front vowels.

2.4   Explaining the trigger/target asymmetries

We have seen from Linker's and Terbeek's studies, as well as Stevens' discussion of the tongue shape-dependent acoustic consequences of lip rounding, that the manifestations of lip rounding are not the same for all rounded vowels. I will argue that these phonetic differences amongst rounded vowels gives rise to the differences exhibited in their phonological patterning.

2.5   What makes a good trigger?

From the preceding discussion, it appears that those vowels for which the addition of lip rounding induces a relatively weak acoustic effect are the typologically preferred triggers of harmony, whereas those for which the acoustic effect of lip rounding is relatively dramatic are the typologically preferred rounding harmony targets. This can be understood under the assumption that vowel harmony is essentially a perceptually driven phenomenon, an approach first put forth by Suomi (1983). Suomi proposes that harmony is best regarded as a means by which to enhance the probability that a given contrast or set of contrasts will be accurately perceived by the hearer. The key idea is that harmony gives rise to an extension of the temporal span associated with some perceptually vulnerable quality, represented below as [±F]. By increasing the listener's exposure to the quality in question, harmony increases the probability that the listener will accurately identify that quality. Suppose that two competing representations for a given string are available, those given in (5a) and (5b):

(5)   a. C V C V C V   b. C V C V C V

[±F] [±F] [±F]   [±F]
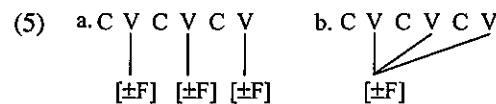
The decision to prefer (b) over (a) has the positive consequence that it provides the listener with increased exposure to the feature value in question. Harmony

gives rise to the perceptual enhancement of the [±F] contrast by extending its duration, although it does so at the cost of reducing the number of possible words in the language.

The harmonic structure in (b) has an additional advantage over the structure in (a). Suppose the listener knows that a given feature is harmonic and thus that over some span the value of that feature will remain constant. Over that span, then, the value of [F] must be identified only once. If the identification is made early on in the string, the acoustic dimension associated with the harmonic feature need no longer be attended to, and attention may be focused on other aspects of the acoustic signal. If only a tentative identification of the harmonic feature value is made early on, additional input is available in the remainder of the string for verification. Finally, if the acoustic cues of the feature in question are somehow obscured in the early portion of the string, the feature value is still potentially recoverable from information carried in the latter portion of the string.

There is a second way in which harmony could be argued to facilitate the correct identification of the triggering vowel. It is well known that vowels exert a coarticulatory effect on neighbouring vowels. Both anticipatory and carry-over coarticulatory effects have been documented for languages such as English (Bell-Berti and Harris 1976), Russian (Purcell 1979), and Catalan (Recasens 1984). In a given V_iCV_j utterance, the articulation of V_i will typically affect that of V_j, and vice versa. It seems reasonable to assume that in VCV utterances in which the vowels are identical or similar, coarticulatory effects will be either nonexistent or fairly minor. If the goal is to maximise the perceptibility of a given vowel, then by insisting that vowels in neighbouring syllables be identical or similar to that vowel, the effects of coarticulation will be eliminated or at least reduced.

Under a perceptual approach to harmony, we would expect that the perceptual advantage of having harmony would be greater when the contrastive value is particularly vulnerable to misidentification. This perceptual motivation underlies the observed trigger preferences. Due to the relative phonetic weakness of rounding in front and nonhigh rounded vowels, harmony triggered by these vowels creates a greater perceptual advantage than harmony triggered by their phonetically more stable high and back counterparts.

To summarise, then, the proposal is that vowel harmony is a perceptually driven phenomenon that serves to prolong the duration of a given feature or quality. As such, its utility is greater when the prolonged quality is one that is particularly at risk for misidentification. Rounding harmony triggered by nonhigh vowels and front vowels (i.e. 'bad' rounded vowels) will thus perform a more critical function than harmony triggered by the perceptually more stable high and back rounded vowels (i.e. the 'good' rounded vowels). Similarly, as I will discuss below in section 2.8, the utility of harmony initiated by a

prosodically short trigger will be greater than that of harmony triggered by a prosodically long trigger.

2.6   What makes a good target?

Rounding harmony is frequently blocked when its application would give rise either to a nonhigh rounded vowel or to a front rounded vowel. One might attribute this pattern to the relative markedness of these vowels, as formalised in Chomsky and Halle's marking convention #XI (1968: 405), and later in Archangeli and Pulleyblank's theory of Grounded Phonology (1994: 78). Under a markedness analysis, the claim would be something to the effect that rounding harmony rules are less highly valued when they generate marked feature combinations. Such an approach relies on the proposition that markedness is a property of individual segments, a view that is challenged by Flemming (1995, this volume), who claims that markedness should be understood as a property of contrasts, not segments.[3] He argues, for instance, that the relative rarity of front rounded vowels and back unrounded vowels within vowel inventories is not due to their inherent markedness. Rather, it is 'the contrast between front rounded and back vowels that is marked because it is less distinct than a contrast between a front unrounded vowel and a back vowel' (Flemming 1995: 26).

Under Flemming's view, there is indeed a connection between markedness and the likelihood that a particular feature will be harmonic within a given language (e.g. rounding in languages in which rounding and backness are independently contrastive, nasality in languages with nasal and oral vowels, and ATR/RTR). Rather than attributing the typological preference for unmarked harmony targets to markedness per se, I believe that a more plausible explanation is one based on the role of perceptual influences in the evolution of rounding harmony systems.

We may assume that when the acoustic consequences of adding rounding to a given tongue shape are strongest, rounding will be most reliably recovered by listeners. Those vowels that are less reliably perceived as rounded might be interpreted as not having undergone harmony, despite the speaker's production of lip rounding during the articulation of these vowels. Over successive generations of speakers, the less robust targets may be reinterpreted as nontargets and a system might emerge in which harmony targets only those vowels whose roundedness has been most reliably perceived. The claim, then, is that the trigger preferences and the target preferences share a common perceptual origin. Harmony triggered by perceptually less salient vowels will perform a greater functional service, and those vowels that are perceptually more salient will be more likely to be interpreted by listeners as having undergone harmony.

2.7   *Why should triggers and targets agree in height?*

I focus now on the avoidance of cross-height harmony, that is the relative rarity of harmony in configurations where the trigger and target disagree in height. I propose that this pattern reflects a requirement that a given articulatory instruction, or autosegment, have a uniform execution mechanism throughout its span of association. In other words, a single autosegment should be interpreted phonetically as an instruction to achieve a single target articulatory posture. Cross-height harmony is thus avoided because the lip rounding gesture is not equivalent for high and nonhigh rounded vowels: typically, high vowels are more rounded than nonhigh, and there are also sometimes effects of backness as well. In this section I will consider some of the phonetic literature that supports this point.

Goldstein (1991), examining data from Linker (1982), concludes that the articulatory goal for rounded vowels is contact along the sides of the upper and lower lips: '[w]hat is specified is whether or not the *upper and lower lips touch along their sides*' (Goldstein 1991: 98). This single factor very clearly separates Linker's lip-activity measurements for rounded vs unrounded vowels. Rounded vowels involve side contact of the lips; unrounded vowels do not. Within the class of rounded vowels, there is a relation between jaw height (and consequently lip aperture) and amount of side contact:

[w]hen the lips are touching there will be an inherent relation between LA [lip aperture] and LW [lip width]. As LA decreases (everything else being equal), the length of the lips' contact region along the sides will increase, and the side-to-side width of the *opening* decreases.   (Goldstein 1991:100)

A harmony span including vowels mismatched for height would necessarily involve re-adjustments in lip aperture and lip width. The avoidance of cross-height harmony can thus be construed as the avoidance of this kind of articulatory readjustments – a principle that I will label Gestural Uniformity.

Note that while there is a clear difference in side contact between high vs nonhigh rounded vowels, there is no systematic difference in length of side contact among front vs back rounded vowels (Goldstein p.c.). This may suggest that a gestural uniformity violation should be assigned to a trigger-target pair of differing heights, but not to a trigger-target pair disagreeing in backness. This is consistent with the typological patterns described here, although the prediction is difficult to test because many rounding harmony languages (the Turkic languages in particular) also exhibit backness harmony.

A putative phonological constraint that refers to the phonetic realisation of a phonologically multiply-linked structure can only be accepted if the geometry of the phonological representation is directly reflected in the phonetic outcome, as in (6a). This approach would be invalidated if one could show that the actual

phonology-to-phonetics mapping is one in which the phonetics reconfigures the phonological output, conferring a rounding target on each vowel as in (6b).

(6)   Possible phonology-to-phonetics mappings
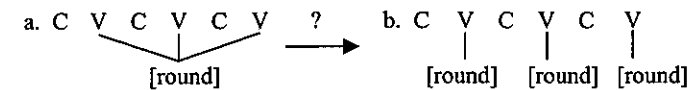
Boyce's (1988) study of coarticulation in English and Turkish provides experimental evidence that bears on this question. Boyce studied vowel-to-vowel coarticulation in English and Turkish uC_0u utterances. These two languages were chosen for comparison because there is good reason to believe that segmentally identical sequences may be assigned distinct phonological representations in these languages. Turkish, as a rounding harmony language, arguably represents uC_0u sequences as containing a single [round] autosegment multiply linked to both vocalic positions. English, which lacks rounding harmony, would be expected to represent the same sequence with two independent [round] specifications:

(7)   Hypothesised phonological representation of uC_0u sequences for English and Turkish
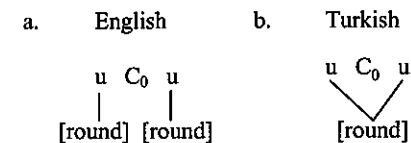
The question investigated by Boyce was whether the distinct phonological representations shown in (7) correspond to distinct articulatory patterns.

The English articulatory pattern, based upon measurements of lip activity and position, yielded a 'trough'-like pattern, shown schematically in (8). The tracing represents lip protrusion (especially of the lower lip):
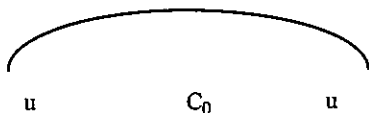
(8)   English 'trough' pattern
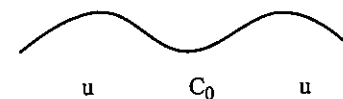
As indicated, the lips attained a position of protrusion in the articulation of the first rounded vowel, then receded during the articulation of the consonantal sequence, then once again attained a position of protrusion for the second rounded vowel.

The Turkish articulatory pattern was qualitatively different. The results obtained by Boyce showed a 'plateau'-like pattern in the articulation of uC_0u sequences by Turkish speakers. This is shown schematically in (9):

(9)   Turkish 'Plateau' pattern

In the Turkish articulation, as shown, the lips attained a position of protrusion during the articulation of the first rounded vowel and remained protruded throughout the utterance.

These experimental findings suggest that whereas the English speakers executed two lip rounding movements, the Turkish speakers executed only one, indicating that the distinct phonological representations appropriate for English and Turkish give rise to distinct phonetic behaviour. If the Turkish pattern is representative of harmony languages in general, then we have reason to reject the remapping in (6b) and conclude that a single [round] autosegment in the phonology corresponds to a single lip rounding gesture in the phonetics. Gestural uniformity regulates the production of such multiply linked structures, dictating that the execution of a single autosegment or articulatory instruction should be achievable by the assumption of a uniform articulatory posture.

2.8   Length-based trigger asymmetries

If the tendency of front and low vowels to trigger harmony has a perceptual origin, then we would expect other factors that impede perception of vowel quality likewise to be involved in the typology of harmony triggers. One such factor, discussed inter alia by Crosswhite (this volume), is length: the quality of shorter vowels is harder to perceive. If a prosodically short vowel is more prone to misidentification than the corresponding long vowel, the perceptual account of harmony would predict that, in cases in which the length of the trigger is relevant to the applicability of harmony, short vowels should be the preferred triggers. In fact, there is evidence that vowel length can play a role in rounding harmony in just this way.

I am aware of three cases in which short vowels trigger harmony while their long counterparts do not. As far as I know, the reverse asymmetry, in which rounding harmony is triggered by long rounded vowels but not by their short counterparts, is unattested.[4] A particularly convincing case of this pattern is the Southern Tungusic language Baiyina Orochen (Li 1996), in which rounding

harmony is triggered only by short vowels. Baiyina Orochen exhibits the general Tungusic pattern of harmony in which vowels within a word agree with respect to tongue root advancement. Alongside this general tongue root harmony exists a more restrictive harmony system involving rounding. Only nonhigh vowels trigger and undergo rounding harmony, also the typical Tungusic pattern. Consider the words in (10), where progressive rounding harmony is triggered by a short vowel. Targets are underlined:

(10)   Short triggers (Li 1996: 126)
  a. tɔyɔ   'fire'   c. tʃolpon   'morning star'
  b. ɔpɔɔ   'rocky hillock'   d. oɲtot   'strange'

Long vowels do not trigger rounding harmony, as shown in (11):

(11)   Long nontriggers (Li 1996: 126)
  a. kɔɔxan   'child'   c. koorgə   'bridge'
  b. ɔɔŋan   'mountain pass'   d. oodən   'velvet'

It is important to note that while long vowels do not *initiate* a rounding harmony domain, they may both undergo and, more importantly, *propagate* harmony. This is shown in (12):

(12)   Long vowels propagate harmony (Li 1996: 131)
  a. gɔlɔɔ-tkɔɔɪ   'log; direct.'
  b. ɔpɔɔ-lɔɔ   'rocky hillock; destin.'
  c. sokkoo-mɲo   'muddy (water); contem.'
  d. oloo-no-tʃo-   'to cook; intent. asp.; pt.t.'

This pattern is consistent with the perceptual account of rounding harmony outlined above. If the functional advantage of harmony is to increase the span of a particular distinctive (and hence important) quality, then the fact that only short vowels initiate harmony can be attributed to their relative perceptual vulnerability. The fact that both long and short vowels may serve to propagate the harmonic feature reflects the fact that these vowels do not carry distinctive information for the feature in question, hence no length-based asymmetry is expected in their harmonic behaviour.

A similar pattern is observed in the Northern Tungusic language Evenki (Nedjalkov 1998). In Evenki, only short vowels trigger rounding harmony; long vowels do not. So when the accusative definite morpheme /–vA/ is added to 'fish', the resulting form is *goro-vo* with a rounded suffix vowel. When this suffix follows a long vowel, rounding harmony does not occur, as in the affixed form of 'tree', realised as *moo-va*, rather than **moo-vo*. Additionally, Harrison (1999a) presents a case from Tuvan, a Turkic language of Siberia, in which short vowels trigger rounding harmony while their long counterparts do not.

2.9   *Summary*

I have argued in this section that harmony is fundamentally a perceptually driven phenomenon that serves to prolong the duration of some contrastive quality. To summarise:

- Acoustic, articulatory, and perceptual data were introduced to support the claim that contrastive rounding is particularly subtle for both nonhigh and front vowels. These are the vowels that are typologically preferred as rounding harmony triggers.
- Nonhigh and front vowels are typologically *dispreferred* as rounding harmony targets. This dispreference was attributed to the greater salience of rounding in the preferred targets which, over time, might lead to the retention of only those vowels as legitimate grammatical targets of harmony.
- The avoidance of cross-height harmony was linked to the difference in lip postures associated with high vs nonhigh rounded vowels, by means of the proposed gestural uniformity constraint.

3   A constraint-based account

It is simple enough to characterise each of the attested rounding harmony systems within a linear, rule-based framework such as that developed in Chomsky and Halle 1968, or within an autosegmental framework such as that employed in Clements and Sezer 1982. For instance, for Type 5 languages one could posit the rule in (13), which indicates that a high vowel assimilates in rounding to a preceding rounded vowel.
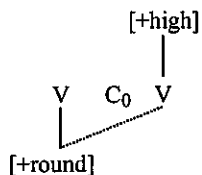
(13) SPE-style rule for Type 5

$$\begin{bmatrix} +syl \\ +high \end{bmatrix} \rightarrow [+round]/ \begin{bmatrix} +syl \\ +round \end{bmatrix} C_0 \text{---}$$

A high vowel is rounded following a rounded vowel.

Alternatively, one might represent the vocalic features autosegmentally, and posit a rule such as that in (14):

(14) Autosegmental rule for Type 5

Rounding spreads from a rounded vowel onto a following high vowel.

Similarly, Type 2 could be represented with the rule in (15) (or some autosegmental analogue of it):

$$(15) \begin{bmatrix} +syl \\ \alpha high \end{bmatrix} \rightarrow [+round]/ \begin{bmatrix} +syl \\ \alpha high \\ +round \end{bmatrix} C_0 \text{---}$$

A vowel is rounded if it is preceded by a rounded vowel of the same height.

These frameworks are capable of describing all of the attested rounding harmony patterns. However, given the mechanisms made available by these two formal systems, there is no way of distinguishing the attested rounding harmony rules from many formally similar but typologically unattested ones. For instance, the linear rule in (16) is no less complex than those in (13) and (15), yet it is not known to play a role in the grammar of any language.

(16) Formally similar rules

$$\begin{bmatrix} +syl \\ -high \end{bmatrix} \rightarrow [+round]/ \begin{bmatrix} +syl \\ +round \end{bmatrix} C_0 \text{---}$$

A nonhigh vowel is rounded when it is preceded by a rounded vowel.

The same point can be made with regard to the formal mechanisms of the autosegmental model. Thus, the rule-based approaches in no way limit the range of predicted rounding harmony systems.

More problematic for rule-based accounts is the fact that the formally most simple rounding harmony rule – one in which a rounded vowel triggers rounding of a neighbouring vowel without regard to height or backness – is typologically quite rare. If a maximally simple rule is assumed to be more highly valued than a more complex rule, then these models make the incorrect prediction that the harmony system characterised by the rules in (17) – i.e. Type 1 harmony – should be the most widely attested.

(17) Formally simplest, but typologically very rare rounding harmony rules

$$[+syl] \rightarrow [+round]/ \begin{bmatrix} +syl \\ +round \end{bmatrix} C_0 \text{---}$$

A vowel is rounded when preceded by a rounded vowel.

Optimality Theory (hereafter OT) replaces the rules of earlier generative models with constraints. The constraints are understood to be supplied by Universal Grammar (UG), while their relative importance or ranking is determined on a language-specific basis. OT is thus implicitly a model of linguistic typology in that a possible grammar is any ranking of the (fixed) set of universal constraints. A model such as this should allow for the characterisation of the conditions from (2) in the form of explicit grammatical statements or constraints.

3.1   Constraints

Following Smolensky (1993), I will assume that the grammatical expression of harmony is by means of Alignment (McCarthy and Prince 1993). Alignment constraints call for the coordination of domain edges. In the case of rounding harmony, the phonological domain, or feature, will be [round], while the morphological domain with which it is aligned will typically be the Prosodic Word (though see the discussion of Sibe in section 3.3).

(18) The pro-harmony constraint: ALIGN

ALIGN-L/R([RD], PRWD)   *The autosegment [round] is aligned with the L/R edge of the Prosodic Word.*

One violation of this constraint is assessed for each docking site (vowel) following the last docking to which a [round] autosegment is linked.

I will also assume two further constraints, whose empirical effects tend to involve the selection of specific triggers for harmony. The alignment constraints in (19) and (20) each introduce a particular refinement to that given in (18):

(19) Nonhigh trigger

ALIGN-L/R ([RD/–HI], PRWD)   *The autosegment [round], when co-occurring with [–high], is aligned with the L/R edge of the Prosodic Word.*

(20) Front trigger

ALIGN-L/R ([RD/–BA], PRWD)   *The autosegment [round], when co-occurring with [–back] is aligned with the L/R edge of the Prosodic Word.*

These more specific alignment constraints, if ranked above general [round]-alignment, will give the effect of promoting harmony when the potential trigger is one of the preferred trigger-types, either a front rounded vowel or a nonhigh rounded vowel. Thus, the formal account of the observed trigger preferences lies in the absence in UG of constraints specifically promoting harmony triggered by the perceptually more salient high and back vowels.

The grammatical mechanism that suppresses harmony, that is, the constraint that overrides the alignment constraints in languages lacking rounding harmony, will be represented by means of a Faithfulness constraint from the DEP family (McCarthy and Prince 1995):

(21) The anti-harmony constraint: DEP

DEP(LINK)   *The output may contain no association line absent in the input.*

To account for the vast majority of vowel harmony systems, in which harmony fails to apply in certain contexts, the grammar must also contain constraints whose effect is to block harmony when it would apply to one of the dispreferred target configurations.

First, I assume a general constraint, labelled *RoLo, which states a dispreference for nonhigh rounded vowels, thus constituting the grammatical manifestation of the articulatory and perceptual bias against lower rounded vowels discussed in section 2. This constraint, or some related grammatical principle, is instrumental in shaping vowel inventories, which commonly lack low rounded vowels. In harmony, *RoLo serves to block harmony when the output would contain a nonhigh rounded vowel not present in the input, that is, harmony targeting nonhigh vowels:[5]

(22) *RoLo[6]   *Nonhigh rounded vowels are avoided.*

In addition, also in keeping with cross-linguistic vowel inventory patterns, I will posit a similar constraint against front rounded vowels. This constraint, like *RoLo, reflects the articulatory and perceptual bias against front rounded vowels:

(23) *RoFro   *Front rounded vowels are avoided.*

Finally, we need a constraint forcing vowels within a rounding harmony span to share the same height specification. I argued above that this reflects a phonetic imperative to avoid the need for articulatory adjustments in the execution of a single gesture. This preference presumably has as its grammatical expression a family of gestural uniformity constraints. In rounding harmony systems, gestural uniformity will reject any instance of the autosegment [round] linked to positions with distinct height:

(24) Gestural uniformity

GESTUNI [(ROUND)]   *A multiply-linked [round] autosegment corresponds to a uniform mechanism for the execution of [round].*

The constraint inventory just laid out suffices to characterise all of the rounding harmony types identified in section 2. Space does not permit the full analysis to be given, but the crucial ranking needed to generate all thirteen types are listed in table 4.5. To illustrate how the constraints interact, I will present an account of one of the more complex languages surveyed, namely Sibe. This language brings together most of the relevant phenomena into a single example.

Table 4.5. *A hierarchical representation of the typology*

Type 1	ALIGN[RD] » *others*
Type 2	UNI[RD] » ALIGN[RD] » *others*
Type 3	UNI[RD], *ROLO » ALIGN[RD] » *others*
Type 4	UNI[RD] » ALIGN[RD/–HI] » DEP(LINK) » *others*
Type 5	*ROLO » ALIGN[RD] » *others*
Type 6	ALIGN[RD/–HI] » UNI[RD] » ALIGN[RD] » *others*
Type 7	ALIGN[RD/–BA], ALIGN[RD/–HI] » UNI[RD] » ALIGN[RD] » *others*
Type 8	ALIGN[RD/–BA] » UNI[RD] » ALIGN[RD] » *others*
Type 9	ALIGN[RD/–BA] » UNI[RD], *ROLO » ALIGN[RD] » *others*
Type 10	*ROFRO » ALIGN[RD/–HI] » *ROLO » ALIGN[RD] » *others*
Type 11	*ROFRO, *ROLO » ALIGN[RD] » *others*
Type 12	*ROFRO » UNI[RD] » ALIGN[RD/–HI] » DEP(LINK) » *others*
Type 13	*ROLO » ALIGN[RD/–BA] » DEP(LINK) » *others*

3.2   Sibe

In Sibe, a south-west Tungusic language described in Li 1996, a rounding harmony pattern similar to that observed in Type 6 is exhibited within roots: as with other Type 6 languages, cross-height harmony is tolerated only with a preferred (i.e. high) target. Nonhigh vowels undergo rounding harmony as long as the trigger is also nonhigh, while high vowels undergo harmony without regard to the height of the trigger. Sibe is more complicated, however, in that while the language has front rounded vowels, only back vowels surface as the product rounding harmony. Sibe thus evidences both target preferences – the preference that targets be high and the preference that targets be back. Examples are shown in (25). In (a–b), the target is high. In (c–d), the target is nonhigh and agrees with the trigger in height. In (e–g), harmony fails to apply either because of height disagreement and a nonhigh target (e), or because the potential target is front (f–g):

(25) Root harmony in Sibe (Li 1996: 195–6)[7]

a.	fulxu̱	'root'
b.	çögu̱	'vegetable'
c.	ɔmɔ̱l	'grandson'
d.	ölχɔ̱	'cowardly'
e.	uva̱ (*uvɔ)	'flour'
f.	utçi̱ (*utçü)	'door'
g.	çöbɛ̱ (*çöbö)	'subsidiary'

Harmony targeting suffixes is even more restricted. As is the case within a root, while high back vowels undergo rounding harmony (26a–c), front vowels in suffixes are never targeted by rounding harmony (26d–e):

(26) Affix harmony in Sibe (Li 1996: 199–204)

a.	batur-lu̱	'to be heroic'
b.	bɔmbɔn-nu̱	'to form a cluster'
c.	gö-χu̱	'to hit (the target)'
d.	bu-ki̱n (*bu-kün)	'to give'
e.	χɔsu-tçi̱ (*χɔsu-tçü)	'corner'

Where suffix harmony differs from root harmony is in the fact that only high suffix vowels undergo harmony. Li's discussion clearly indicates that nonhigh vowels in suffixes never undergo harmony, even if they agree in height with the potential trigger. He does not, however, include any specific examples of a suffixal /a/ following a nonhigh rounded vowel. Some examples of suffixal /a/ are listed in (27). An invented root in (27e) demonstrates the crucial context in which harmony reportedly fails to apply:

(27) Suffixal /a/ Sibe (Li 1996: 199–204)

a.	ɔvu-ma-χι	'to wash-pres.prog.'
b.	vɣili-rtan	'to work-nom.'
c.	suxu-maq	'axe-instr.'
d.	is-maq	'soap-instr.'
e.	ɔs-maq (*ɔs-mɔq)	*invented*-instr.'

The Sibe pattern is instructive in a number of respects. First, different harmony patterns are exhibited within different morphological domains. We saw that nonhigh vowels are targeted by harmony only within roots. This pattern is easily captured by means of alignment, which calls for the co-incidence of phonological and morphological domain edges.

To account for the Sibe system, we may posit the constraint sub-hierarchy in (28):

(28) Constraint sub-hierarchy for Sibe[8]
   *ROFRO » ALIGN([RD/–HI], ROOT) » *ROLO » ALIGN([RD], PRWD) » *others*

The hierarchy in (28) may be interpreted as follows. First, the fact that front rounded vowels never arise as the result of harmony is reflected in the high ranking of *ROFRO. Next, by virtue of its ranking above the constraint that requires harmony within the domain of the prosodic word, *ROLO blocks the application of harmony when it would target a low suffix vowel. The root harmony constraint (ALIGN([RD/–HI], ROOT)) outranks *ROLO, however, thus allowing nonhigh vowels within roots to undergo harmony. Tableaux are shown in (29) and (30) to demonstrate how the constraints interact to characterise the Sibe rounding harmony patterns:

(29) Root-internal harmony in Sibe

	*RoFro	Align([Rd/–hi], Root)	*RoLo	Align([Rd], PrWd)	Uni[Rd]
☞ utçi					
utçü	*!				
fulxı				*!	
☞ fulxu					
çögı		*!			
☞ çögu					
ɔmal		*!			
☞ ɔmɔl					
☞ uva					
uvɔ			*!		

(30) Suffix harmony in Sibe

	*RoFro	Align([Rd/–hi], Root)	*RoLo	Align([Rd], PrWd)	Uni[Rd]
☞ bu-kin					
bu-kün	*!				
batur-lı				*!	
☞ batur-lu					
gö-χı				*!	
☞ gö-χu					
☞ ɔs-maq					
ɔs-mɔq			*!		
☞ suxu-maq					
suxu-mɔq			*!		

3.3   Factorial typology

Using the program OTSoft (Hayes et al. 2000), I generated a factorial typology based on the seven proposed constraints.[9] The input assumed a generic Turkic-style system with a symmetric eight-vowel system and backness harmony. Thus, for each of the four rounded vowels as potential harmony triggers, two trigger-target pairs were considered: one in which the potential target was high, and

one in which the potential target was nonhigh. This yielded eight two-way choices:

(31) Input to the factorial typology

Trigger	Target	Harmonic		Nonharmonic
u	high	u-u	or	u-ı
u	nonhigh	u-o	or	u-a
o	high	o-u	or	o-ı
o	nonhigh	o-o	or	o-a
ü	high	ü-ü	or	ü-i
ü	nonhigh	ü-ö	or	ü-e
ö	high	ö-ü	or	ö-i
ö	nonhigh	ö-ö	or	ö-e

Given eight two-way choices, there were 2^8 possible outcomes, including the case in which no harmonic pairs were selected. There were thus $2^8 - 1$ ($= 255$) possible harmony systems. Of these, the proposed constraint set generated only thirty-six. All of the thirteen observed types from section 1 were generated. The predicted typology is thus reasonably small, and the cases of overgeneration in general look to be accidental gaps, not systematic ones. Thus, for instance, the pattern in which rounding applies to *ul* and *ol* sequences is the same as the familiar Type 5 pattern (cf. Turkish), except that harmony applies only amongst back vowels. Similarly, the pattern in which harmony applies to *ül* and *öA* is essentially the same-height harmony of Type 2, except that only front vowels participate.[10]

3.4   Cases of free variation

In assessing the typological validity of the constraint set proposed here, we also need to address the cases where the harmony pattern exhibits free variation. We have seen three cases here, all in section 1: Tuha, Tofà, and Altai Tuvan. As a model for free variation in Optimality Theory, I adopt the approach of Anttila (1997a, 1997b), under which certain constraints can be critically freely ranked. In such a grammar, the outcomes that arise for all permutations of the free ranked constraints are considered to be possible outputs.

In Tuha, the constraints that must be freely ranked are gestural uniformity and Align[Rd]: where the former dominates the latter, cases like /o I/ emerge as [o i]; whereas under the opposite ranking /o I/ surfaces as [o u]. The case of Tofà can be analysed similarly, with a critical free ranking of Align[Rd/–ba] and Dep(Link).

The case of Altai Tuvan is harder. Using the constraints proposed here, it is possible to achieve a rough statistical match to Harrison's (2001) observations,

adopting the gradientised version of free ranking discussed in Boersma 1997 and Boersma and Hayes 2001. To achieve an exact match requires the addition of a new constraint to the system, which would require alignment for the 'best' possible triggers, namely vowels that are *both* front and low.[11] The typological consequences of such a constraint remain unexplored, however.

4   Optimality Theory as a model of typology or grammar, or both?

One central claim of OT is that all languages share a common set of (universal) constraints, but differ from one another with respect to how those constraints are ranked. This model entails that for any given grammar, a great number of *inactive* constraints must be present – constraints ranked too low to have any decisive effect in the determination of output forms. The evidence for these constraints is purely typological; that is, that together they form an inventory that generates the typology of existing languages. However, it is a fairly strong claim to say that the constraints are actually present in the grammars of languages in which they are inactive.

I conclude this chapter by presenting experimental results suggesting that speakers do in fact carry around inactive constraints, and that speakers recruit such constraints in the evaluation of novel phonological contexts such as those introduced in loanwords. Ross (1996) has presented a similar argument for Tagalog, as have Ringen and Heinämäki (1999) for Finnish. The experiment described here involves loanwords in Turkish.

4.1   Regressive harmony in Turkish loanwords

Yavas (1980) and Clements and Sezer (1982) demonstrate that in addition to the progressive harmony observed in native Turkish words, regressive harmony is exhibited in loanwords. Such harmony targets epenthetic vowels introduced to break up an initial consonant cluster. These epenthetic vowels are always high, but vary in backness and rounding on the basis of the quality of the first full vowel in the word. For example, backness harmony results in an epenthetic [i] in *siteno*, 'steno', in which the first full vowel is front, but an epenthetic [ı] appears in a word like *sıtar* 'star', where the full vowel is back.

The epenthesis phenomenon is of interest here because the epenthetic vowels sometimes undergo regressive rounding harmony, triggered by the stem; thus *flüt* 'flute', is realised as *fülüt*, with a rounded epenthetic vowel. Reports differ as to the conditions under which regressive harmony will apply. In the variety reported by Yavas, regressive harmony may be triggered only by high vowels, whereas in the variety described by Clements and Sezer, rounding harmony is consistently triggered by high vowels but may sometimes also be triggered by nonhigh vowels.

In Kaun 1999, I described the results of an experiment originally designed to resolve this discrepancy. Nine native speakers of Turkish[12] were presented with 107 loanwords taken from Özgüler 1989. Each of these words contained an initial consonant cluster, and the subjects were asked to indicate the appropriate quality (or qualities) of the epenthetic vowel. The task seemed to present little difficulty to the subjects; however, a great deal of subject-to-subject variation was observed.

All subjects agreed with the Yavas and Clements and Sezer patterns in one respect: High rounded vowels consistently triggered rounding harmony, as in words like *buluz* 'blouse',[13] and *fülüt* 'flute'. When the potential trigger was nonhigh, six distinct rounding harmony patterns emerged. These are presented in (32). In this chart, solid lines enclose consistent rounding harmony triggers, while dashed lines enclose optional triggers. Unenclosed vowels never trigger harmony for the pattern in question. The number of subjects instantiating each pattern is indicated in parentheses:

(32) Six rounding harmony patterns

Group A	(1)	ü	u	ö	o
Group B	(1)	ü	u	ö	o
Group C	(2)	ü	u	ö	o
Group D	(1)	ü	u	ö	o
Group E	(2)	ü	u	ö	o
Group F	(2)	ü	u	ö	o

The chart in (33) includes some of the words used in the experiment:

(33) Sample words

	'flute'	'blouse'	'flirt'	'block'
Group A	fülüt	buluz	filört	bılok
Group B	fülüt	buluz	filört/fülört	bılok
Group C	fülüt	buluz	fülört	bılok
Group D	fülüt	buluz	filört/fülört	bılok/bulok
Group E	fülüt	buluz	fülört	bılok/bulok
Group F	fülüt	buluz	fülört	bulok

Two generalisations can be made on the basis of these patterns. First, we note that, while subjects consistently exhibited same-height harmony (*fülüt* and *buluz*), cross-height harmony was not unanimously applied. Second, with harmony involving a height mismatch, front vowels were preferred as harmony triggers, that is, regressive harmony is observed more frequently in *flört*-type words, where the potential trigger is front, than in *blok*-type words, in which a back vowel serves as the potential trigger of regressive rounding harmony.

These patterns are not predictable on the basis of the progressive rounding harmony pattern of native Turkish words. As noted above in section 1, Standard Turkish is essentially a Type 5 language in which harmony may be triggered by any rounded vowel, but targets only high vowels. Thus, in the native pattern of progressive harmony, while vowel height does play a role in determining the applicability of harmony, it is the *target* whose height is relevant, rather than the trigger. Moreover, backness never serves to restrict the applicability of native rounding harmony.

While the observed patterns of rounding harmony share no common features with the native progressive harmony pattern, they do resemble the general cross-linguistic patterns of rounding harmony. In particular, the avoidance of cross-height harmony along with the preference for front rounded triggers are familiar from the general typology of the phenomenon. I argued in Kaun 1999 that the behaviour of the experimental subjects can be modelled as involving the recruitment of constraints that are inactive in native Turkish, but predicted to exist on the basis of the general typology of rounding harmony. These are among the constraints posited in section 3.

4.2   Discussion and conclusions

The results of the Turkish loanword experiment support the choice of OT rather than a rule-based system to account for the typology of rounding harmony. A rule-based account, which might characterise native Turkish harmony with a rule like that shown above in (13), makes no predictions regarding the realisation of rounding harmony in the loanword context. To the extent that such analyses could be said to offer any predictions with respect to regressive rounding harmony in Turkish, they would either predict that rounding harmony should not occur (because rounding harmony is progressive in native Turkish), or that rounding harmony will always proceed as in pattern F (example (32)), where any rounded vowel triggers rounding of the high epenthetic vowel. Neither of these predictions characterises the observed facts.

The loanword data also indicate that while OT allows for the characterisation of typologies in a direct and falsifiable manner (i.e. an analysis is wrong if a language can be shown to exhibit a pattern that cannot be generated by some ranking of the proposed constraint set), it also appears to be an appropriate model of individual grammars. The Turkish facts support the claim that constraints that never play a decisive role in determining surface structure are nonetheless present in grammars. Alternatively, these results could be interpreted as indicating that, when confronted with a novel phonological configuration, speakers can invent constraints 'on the spot', and that when they do so, they are guided by the same phonetic pressures that govern grammatical systems in general,

in this case, the phonetic principles that underlie the phenomenon of rounding harmony.

I have argued that the account proposed here is functionally grounded. The functional underpinning of the typology can be held up as a means of understanding the evolution of rounding harmony systems. The Turkish loanword pattern further suggests that, in our effort to understand and model phonological systems, we should look upon substantive principles not just as a means of explaining the recurrence of phonological patterns post hoc, but as fundamental components of grammar, accessed and deployed by speakers of human language.

References

Anderson, Gregory and K. David Harrison (2000). *Tyvan*. Munich: Lincom Europa.

Anttila, Arto (1997a). *Variation in Finnish phonology and morphology*. PhD dissertation, Stanford University.

(1997b). Deriving variation from grammar: A study of Finnish genitives. In Frans Hinskens, Roeland van Hout, and Leo Wetzels (eds.), *Variation, change and phonological theory*. Amsterdam: John Benjamins.

Archangeli, Diana and Douglas Pulleyblank (1994). *Grounded phonology*. Cambridge: MIT Press.

Bell-Berti, Fredericka and Katherine S. Harris (1976). Some aspects of coarticulation. *Haskins Laboratories Status Report on Speech Research* SR-45/46: 197–204.

Boersma, Paul. (1997). How we learn variation, optionality, and probability. *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* 21: 43–58.

Boersma Paul and Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *LI* 32: 45–86.

Bogoroditskij, V. A. (1953). *Vvedenie v tatarskoe yazykoznanie* (Introduction to Tatar linguistics). Kazan: Tatgosizdat.

Casali, Roderic (1993). Cross-height vowel harmony in Ghanaian. Ms, University of California, Los Angeles.

Chomsky, Noam and Morris Halle (1968). *The sound pattern of English*. New York: Harper and Row.

Clements, George N. and Engin Sezer (1982). Vowel and consonant disharmony in Turkish. In Harry van der Hulst and Norval Smith (eds.), *Structure of phonological representations (part II)*. Dordrecht: Foris Publications, 213–55.

Comrie, Bernard (1981). *Languages of the Soviet Union*. Cambridge: Cambridge University Press.

Crothers, John (1978). Typology and universals of vowel systems. In Joseph H. Greenberg et al. (eds.), *Universals of human language*. Stanford: Stanford University Press, vol. 2, 93–152.

Derbyshire, Desmond (1985). *Hixkaryana and linguistic typology*. Austin, Tx: SIL publications and University of Texas, Austin.

Flemming, Edward (1995). *Auditory representations in phonology*. PhD dissertation, University of California, Los Angeles.

Goldstein, Louis (1991). Lip rounding as side contact. In *Actes du XIIème congrès international des sciences phonétiques (Vol I/5: Sessions plénières et semi-plénières)*. Aix-en-Provence: Université de Provence, vol. 1, 97–101.

Harrison, K. David (1999a). Harmony and underspecification in Tuvan. In *MIT Working Papers in Linguistics: Proceedings of SCIL XI*. Cambridge, MA: MIT.

(1999b). Tuvan reduplication and harmony. *BLS 25: Parasession on Caucasian, Dravidian and Turkic Linguistics*.

Harshman, Richard (1970). Foundations of the PARAFAC procedure: Models and conditions for an 'explanatory' multi-modal factor analysis. *UCLA Working Papers in Phonetics* 16.

Harshman, Richard and S. Berenbaum (1980). Basic concepts underlying the PARAFAC-CANDECOMP three-way factor analysis model and its application to longitudinal data. In D. H. Eichorn et al. (eds.), *Present and past in middle life*. New York: Academic Press, 435–59.

Harshman, Richard, Peter Ladefoged, and Louis Goldstein (1977). Factor analysis of tongue shapes. *JASA* 62: 693–707.

Hayes, Bruce, Bruce Tesar, and Kie Zuraw (2000). 'OTSoft', software package, http://www.linguistics.ucla.edu/people/hayes/otsoft/.

Herbert, Raymond and Nicholas Poppe (1963). *Kirghiz manual*. Bloomington: Indiana University.

Homer, Molly Kathryn (1998). *The role of contrast in nasal harmony*. PhD dissertation, *University of Illinois*.

Hsu, Chai-Shune (1993). Tsou phonology. Ms, University of California, Los Angeles.

Kaun, Abigail R. (1994). An Optimality Theoretic account of rounding harmony typology. *WCCFL* 13: 78–92.

(1995). *The typology of rounding harmony: An optimality theoretic approach*. PhD dissertation, University of California, Los Angeles.

(1999). Epenthesis-driven harmony in Turkish. *BLS 25: Parasession on Caucasian, Dravidian and Turkic Linguistics*.

Korn, David (1969). Types of labial vowel harmony in the Turkic languages. *Anthropological Linguistics* 11: 98–106.

Li, Bing (1996). *Tungusic vowel harmony: Description and analysis*. The Hague: Holland Academic Graphics.

Linker, Wendy (1982). *Articulatory and acoustic correlates of labial activity in vowels: A cross-linguistic study*. PhD dissertation, University of California, Los Angeles.

Maddieson, Ian (1984). *Patterns of sounds*. Cambridge: Cambridge University Press.

McCarthy, John (1979). *Formal problems in Semitic phonology and morphology*. PhD dissertation, MIT.

McCarthy John and Alan Prince (1993). *Generalized alignment*. Ms., University of Massachusetts and Rutgers University. Rutgers Optimality Archive 7, http://ruccs.rutgers.edu/roa.html.

(1995). Faithfulness and reduplicative identity. In Jill Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk (eds.), *University of Massachusetts occasional papers in linguistics 18: Papers in Optimality Theory*. Amherst: GLSA, 249–384.

Nedjalkov, Igor (1998). *Evenki*. Cambridge: Cambridge University Press.

Özgüler, Ali (1989). *Türçedemizdeki yabanci kelimeler sözlük*. N.P.: Sim Matbaacılık.

Prince, Alan and Paul Smolensky (1993). *Optimality Theory: Constraint interaction in generative grammar*. Rutgers Optimality Archive 537, http://roa.rutgers.edu/.

Purcell, Edward T. (1979). Formant frequency patterns in Russian VCV utterances. *JASA* 66: 1691–702.

Recasens, Daniel (1984). Vowel-to-vowel coarticulation in Catalan VCV sequences. *JASA* 76: 1624–35.

Ringen, C. and Ovokki Heinämäki (1999). Variation in Finnish vowel harmony: An OT account. *NLLT* 17: 303–37.

Ross, Kie (1996). Floating phonotactics: Variability in infixation and reduplication of Tagalog loanwords. MA thesis, University of California, Los Angeles.

Smolensky, Paul (1993). Harmony, markedness, and phonological activity. Paper delivered at the Rutgers Optimality Workshop 1.

(1995). On the internal structure of the constraint component *CON* of UG. Rutgers Optimality Archive 86, htt://roa.rutgers.edu.

Steriade, Donca (1995). Underspecification and Markedness. In John. A. Goldsmith (ed.), *The handbook of phonological theory*. Oxford: Blackwell, 114–74.

Stevens, Kenneth N. (1998). *Acoustic phonetics*. Cambridge, MA: MIT Press.

Suomi, Kari (1983). Palatal harmony: A perceptually motivated phenomenon? *Nordic Journal of Linguistics* 6: 1–35.

Svantesson, Jan-Olaf (1985). Vowel harmony shift in Mongolian. *Lingua* 67: 283–327.

Terbeek, Dale (1977). *A cross-language multidimensional scaling study of vowel perception*. PhD dissertation, University of California, Los Angeles.

Yavas, Mehmet (1980). *Borrowing and its implications for Turkish phonology*. PhD dissertation, University of Kansas.

Notes

1. For each language, data from eight male subjects were obtained. Photographs were taken simultaneously with audio recordings.

2. Back vs nonback (1) separated the front and back vowels with one exception. On the basis of this factor, the phonetically back vowel [ɪ] was grouped with the front vowel cluster. The vowel [ɚ] fell between the back and front vowel clusters. Back vs nonback (2) grouped the vowel [ɪ] with the cluster of back vowels and placed [ɚ] at the low end of the scale along with {ö, i, e, and ü}.

3. This general notion of contrast is applied to the phenomenon of nasal harmony in Homer 1998.

4. In Maltese (McCarthy 1979), backness/rounding harmony targets short vowels, but not long vowels. One explanation for this pattern might be the fact that short vowels are more subject to gestural overlap than longer vowels. The historical scenario would therefore be that at an earlier stage of the language, no harmony existed. Over time, due to coarticulation from a preceding rounded vowel, the short vowels came to be perceived as phonologically rounded, and a system of rounding/backness harmony entered the grammar.

5. The languages in which nonhigh rounded vowels may not undergo rounding harmony all possess nonhigh rounded vowels on the surface. What they do not tolerate is the occurrence of nonhigh rounded vowels as the output of rounding harmony. In these systems, we may assume that the retention of nonhigh rounded vowels that are not the product of the harmony is insured by an input-output Faithfulness constraint that outranks *RoLo. The same account will apply to those languages

that allow front rounded vowels on the surface, but do not allow their appearance as the product of harmony.

6. Both *RoLo and *RoFro are equivalent to Archangeli and Pulleyblank's (1994: 78) grounded path conditions on the feature [+round] and features referring to height and backness. In particular, *RoLo should be functionally equivalent to the LO/RD Condition and the RD/LO Condition, while *RoFro should do the work of the RD/BK Condition and the FR/RD Condition.

7. Disharmonic *uCɩ* sequences occur, but only in fifteen roots from Old Manchu. Relatively uneducated Sibe speakers are reported by Li to be unfamiliar with these roots (Li 1996: 195).

8. The grammars here omit constraints that ensure that the lexical specifications for rounding in stem vowels are respected. Just what these constraints are probably depends on the language. For languages in which the principles of harmony are respected even within stems, a constraint of the type IDENT([RD], INITIAL SYL- LABLES) is probably appropriate. For languages like Standard Turkish, in which the lexicon is very rich in disharmonic roots, probably the correct constraint is IDENT([RD], ROOT). In either case, the relevant constraint is undominated, so I will simply omit candidates from consideration that alter the rounding of the trigger.

9. I wish to thank Bruce Hayes and Kie Zuraw for help in generating and interpreting the factorial typology described here.

10. The full set of plausible-but-unattested patterns is as follows. In each pattern, {} encloses the vowel pairs that undergo harmony: (1) {uI, öA, oA}, (2) {uI}, (3) {uI, oI}, (4) {uI, oI, uA, oA}, (5) {oI, oA}, (6) {oI}, (7) {öI, uI, oI, öA, oA}, (8) {öI, uI, oI, öA, uA, oA}, (9) {öI, uI, oI}, (10) {öI, oI, öA, oA}, (11) {öI, oI}, (12) {üI, uI, öA}, (13) {üI, öA, oA}, (14) {üI, öA}, (15) {üI}, (16) {üI, öI, uI}, (17) {üI, öI, uI, üA, öA, oA}, (18) {üI, öI, üA, öA, oA}, (19) {üI, öI, üA, öA}, (20) {üI, öI, oI, öA, oA}, (21) {üI, öI, oI}, (22) {üI, öI, oI, üA, öA, oA}, (23) {üI, öI, oI, üA, öA}.

11. On 'conjoined' constraints of this type, see Smolensky 1995.

12. The subjects ranged in age from eighteen to thirty-five years. All resided in the New Haven, CT area at the time of the experiment, but had been raised in urban settings including Ankara, Istanbul, and Izmir.

13. Due to a fronting effect of borrowed [l], this word is transcribed as *bül'uz* in both Yavas 1980 and Clements and Sezer 1982. The subjects in the experiment reported here did not produce front vowels before [l], supplying *buluz* instead.

5   The evolution of metathesis

Juliette Blevins and Andrew Garrett

1   Introduction

Our object of study in this chapter is metathesis, which we define as any re- ordering of segments or features within the phonological string.[1] Representative cases, discussed in more detail below, are illustrated in (1).

(1) a. Rotuman: /mofa/ → *moaf* 'rubbish'; cf. (20) below
        Proto-Indo-European > Armenian: *k̑ubʰros > surb* 'holy' (initial *k̑ > s*); cf. (6) below

     b. Nxilxcín (Colville): *sʕáy* 'they are noisy' vs *sy-m-əncʕàt* 'they make noise'; cf. (15) below
        Classical > South Italian Greek: *gambrós > grambó* 'son-in-law'; cf. (10) below

     c. Marathi: *ōtʰ > hõʈ* 'lip'; cf. (17) below

Adjacent segments seem to exchange positions in the common pattern seen in (1a), while (1b) shows examples of nonlocal movement. A case of feature metathesis is shown in (1c); such cases are relevant because feature and segment metatheses differ in their phonological effects but not their underlying causes.

Metathesis has long posed problems for phonological theory. These problems are of two main types: metathesis has resisted analysis in terms of phonetically natural or motivated sound change, and the reordering of sounds in metathesis has required extensions of otherwise highly restrictive phonological formalisms. We will argue here that metathesis can, despite these problems, be explained in a phonetically natural way based on precisely the same assumptions required to understand other phonological phenomena.

We also have a more programmatic goal. In recent years, phonologists have increasingly come to accept the view that phonological patterns, both within and across languages, can be explained by reference to the findings of experimental phonetics. As yet, however, there is no consensus as to the precise explanatory nexus between the two areas. In this chapter we will contrast two views of the relationship between phonetics and phonology, for which we will use the short-hand terms *phonetic optimisation* and *evolutionary phonology*. The first

Phonetically Based Phonology

edited by

Bruce Hayes, Robert Kirchner, and Donca Steriade

Contents